WSRL-TM-18/91

AR-006-707

AUSTRALIA

A DECISION SUPPORT SYSTEM FOR MILITARY INTELLIGENCE ASSESSMENT

N.G. GORI and P.F. CALDER

COMBAT SYSTEMS DIVISION
WEAPONS SYSTEMS RESEARCH LABORATORY

Approved for Public Release.

MARCH 1991

DSTO WSRL DEPARTMENT OF DEFENCE
DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION

UNCLASSIFIED

DSTO AUSTRALIA

TECHNICAL MEMORANDUM
WSRL-TM-18/91

DTIC QUALITY INSPECTED 5

A DECISION SUPPORT SYSTEM FOR MILITARY INTELLIGENCE ASSESSMENT

Accesion For		
NTIS CRA&I	✓	
DTIC TAB	☐	
Unannounced	☐	
Justification		
By		
Distribution /		
Availability Codes		
Dist	Avail and / or Special	
A-1		

N.G. Gori and P.F. Calder

SUMMARY (U)

This paper details the early stages of research towards designing a decision aid to assist an Intelligence Officer in a tactical HQ. The particular problem being tackled, the interpretation of intelligence messages concerning opposing force activities, is described, as are those factors contributing to the task's difficulty.

Author's address:
Combat Systems Division
Weapons Systems Research Laboratory
PO Box 1700, Salisbury
South Australia

Requests to: Chief, Combat Systems Division

UNCLASSIFIED

TABLE OF CONTENTS

LIST OF FIGURES

1. INTRODUCTION

The use of computers in command and control systems is currently limited to supporting the mechanics of information handling, such as storage, retrieval, collation and display. Computers have been unable to "process" information; drawing conclusions from the available information has remained solely within the province of a human operator. Military intelligence processing is an area of command and control where these limitations are severely evident.

Progress in artificial intelligence suggests that it should be possible to use computers to aid the intelligence process much more effectively (reference 5; Spain 1983). Such a system would draw conclusions based on its knowledge of an opposing force and terrain.

We are in the process of identifying the knowledge required to define an opposing force and the interaction needed with terrain, so that appropriate conclusions can be made when new information is received.

This paper details the early stages of research towards developing a decision aid to assist an Intelligence Officer in a tactical HQ. This decision aid would augment well established technologies with specific artificial intelligence support. The particular problem being tackled, the interpretation of intelligence messages concerning opposing force activities, is described, as are those factors contributing to the task's difficulty.

2. THE MESSAGE EVALUATION PROBLEM

In Combat Systems Division we are interested in problems of Command, Control and Intelligence in the tactical area. We are particularly interested in developing tools to assist an Intelligence Officer in a tactical HQ with providing timely and accurate assessments to his Commander.

From the reported observed actions of enemy forces, (received in the form of written and spoken intelligence messages), the requirement is to derive as much reliable information about opposing forces as is practicable. In particular, the need is to identify enemy units and their location, and to postulate the identities and locations of units not previously detected. We need to assign roles and activities to units, infer their aims, their intentions, and their most likely course of action. We need also to infer their organisational relationships.

2.1 The quality of received information

In an ideal world, the information an Intelligence Officer receives would be of a consistently high accuracy and reliability. Similarly, in an ideal world, he would have sufficient time to accurately assess every message. In the real world, this is patently not the case.

Messages will vary in their quality, both because of the reliability of the source, and because of the methods with which the information was obtained. Both these factors take time to assess.

Generally, every message will define some activity associated with a unit of an opposing force. The Intelligence Officer must decide the identity of the unit, what it is doing, and how this information matches what he would expect.

Received information will invariably be incomplete. For example, a sighting may report only some of the vehicles of a unit, leading to the conclusion of a smaller force than is actually there.

Messages will vary in content. For example, a Direction Finding report may give a rough location, yet provide a good identification of a unit's designation. In contrast, an aerial reconnaissance report may give a good location, but no identification.

It is highly likely that the time available to gather essential information will be very short. When the information eventually arrives, the Intelligence Officer will usually have to assess it quickly. This gives rise to the possibility of overlooking critical correlations.

2.2 Information processing

The Intelligence Officer must maintain a mental model of the opposing force and its activities. He uses this model to define a picture of where the opposing force is and what it is doing.

Each message must be assessed relative to the current picture about the opposing force to try to fit the new information into the pattern. If it matches expectations, confirming some already known data, it increases the probability that the previously reported information is true.

If the new information doesn't match expectations, then either the picture from which the expectations are derived is incorrect, or the information is incorrect. If the validity of the new information is verified from other sources, then the picture must be modified to accommodate this new information.

3. THE PROPOSED SYSTEM

It is accepted that Artificial Intelligence (AI) techniques are necessary to significantly enhance the capabilities of the next generation of command and control systems. But these techniques will not be sufficient on their own. A working system will need to use other software technologies, and will need to interact with other, non-AI systems. A major focus of this work, therefore, is in developing and integrating a number of diverse technologies.

The proposed system will be event driven, where an event is defined as the receipt of new information concerning the opposing force. The system will react to this information and prompt the user with new conclusions.

In this section, we discuss the relevance of component technologies of the proposed system. We also discuss some of the philosophical issues underpinning our approach.

3.1 Visualisation

A key aspect of the work is the necessity for the Intelligence Officer to be an integral part of a working system, summed up in the phrase "having the man in the loop". There are pragmatic and philosophical reasons for this approach.

Being pragmatic, it seems unlikely that a fully autonomous expert system can be built for this domain in the foreseeable future.

The consultant paradigm for human - expert system interaction is not generally successful. A combination of human user and expert system is more likely to cope successfully with planned and unanticipated combinations of events, if the user is able to draw upon his knowledge and experience(ref.12). One must take advantage of human abilities for the system to have the remotest chance of success. The combined capabilities of man and machine are likely to result in a much more robust system, available sooner, than would otherwise be possible.

It would be undesirable to create a system where a user's role is reduced to simply being a passive data gatherer, answering system generated questions without knowing why they are important. In life and death situations, it is unreasonable to expect a user to blindly accept machine generated conclusions, where there are not exceptionally lucid explanations.

For these reasons, the emphasis is on providing advanced visualisation to facilitate complex information exchanges between man and machine. A user will then be more able and willing to guide and influence the system behaviour, and to accept subsequently generated conclusions. The relationships between man, visualisation and inferencing systems is shown in the Visualisation Triangle(ref.14) in figure 1.

The crucial point is that a system must support the cognitive demands and user's representation for a task. The aim is to build representational aids, to create and manipulate representations of the target world(ref.16).

We intend to achieve this high level of visualisation by presenting information in the form that the user is most familiar with. For example, geographical objects will be portrayed relative to terrain against a map background. Diverse information will be displayed to a user as he progresses towards his conclusions. Conclusions generated by the system will provide explanations, both textual and graphical, to justify the chain of reasoning supporting them. Interactions between the user and system will be appropriate to the form of the information displayed.

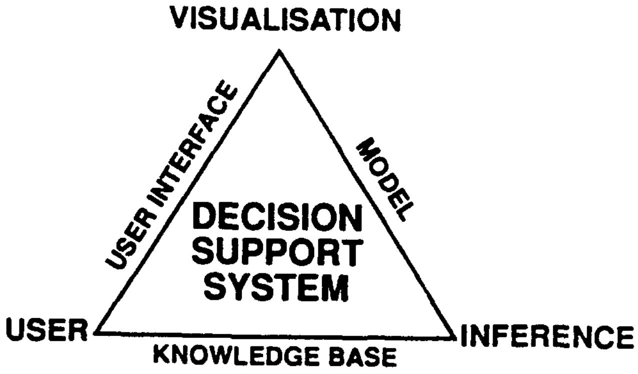

Figure 1. The Visualisation Triangle

3.2 Technologies

3.2.1 Object oriented databases

A cornerstone of this work is the Spatial Object Toolkit (SPOT), which is being developed at the CSIRO Centre for Spatial Information Systems. SPOT will provide a basic terrain data manipulation facility.

SPOT, in turn, is built around the ONTOS object-oriented database system(ref.1). Using an object-oriented database will allow the representation of much more complex data structures than is possible with relational database technology, with the promise of little or no performance degradation(ref.3).

The objects to be stored will include terrain objects, and details of the opposing force organisation. In military terminology, the latter is called the Order of Battle (ORBAT).

3.2.2 Graphical User Interface tools

The Graphical User Interface (GUI) is intended to provide "user-friendly" front-ends, and a common "look and feel" to the different parts of the system. It will be a WIMP system, with the map background forming one of the major windows.

The system will provide facilities to superimpose the tactical situation on a map background. This situation will be represented by scalable linear, area or point objects representing tactical features and terrain objects of military significance, and non-scalable icons representing military units. These will represent stored database objects which will have been created either interactively by the user directly from the screen, iconic features selected from the ORBAT, or the result of system generated conclusions. It is a long term goal to automate large parts of the terrain analysis process.

One of the modules will be an ORBAT manipulation tool. The tool will allow a user to generate an actual ORBAT, or grouping, by selecting and dragging ORBAT items from the doctrinal ORBAT as units are identified. As the system is developed, the reasoning system will be able to infer possible unit identities which match reported sightings.

Another tool would allow the generation of database queries by using menus and pointing to displayed graphical objects.

A necessary tool will be a Message Input module, through which new information on the opposing force enters the system.

3.2.3 *Hypertext system*

The hypertext system should fulfil a number of general aims. It should make easily accessible to a user the copious amounts of diverse information required for an intelligence assessment system. These information sources might include pictures, databases, equipment summaries, ORBATs, and planning tools. Together with the GUI tools, the intention is to combine these various information sources seamlessly, and to minimise the amount of system specific training required.

3.2.4 *Digital map data*

The SPOT system will use digital terrain maps (DTM) containing relief, vegetation, cultural and drainage data stored as topological constructs. Storing the DTM data topologically allows complex terrain objects, such as river or road systems, to be accessed as single objects. It would enable the system to reason about terrain features, because the characteristics of that object would be uniquely identifiable.

It has been necessary to commission the generation of DTM data for this work, as none was available from the usual mapping sources. There remains a great deal of debate as to the optimum representation of terrain data for military command support systems. Part of this research will investigate how well topological representations make terrain data accessible to a computer reasoning system, instead of being displayed solely for human reasoning.

3.2.5 *Artificial intelligence*

So far, this paper has described features which, when implemented, would allow an Intelligence Officer to execute the same tasks he currently performs, in broadly the same way, albeit on a computer based system. The belief is that the minimum benefit to be gained by incorporating AI

technologies is that they would relieve the Intelligence Officer of some of his hack workload, allowing him to spend his time more creatively, and making him less likely to be overwhelmed at times of peak activity.

Much of the work in this project has been in knowledge elicitation and in building up an understanding of the domain. Part of this domain knowledge is an indication of where an automated reasoning system can aid an Intelligence Officer, that is, what aspects are, or might soon become, amenable to machine reasoning, and what aspects are best left to the human in the loop. What is needed is a symbiosis between man and computer(ref.15). We are now in a good position to describe some of the features of the intelligence assessment problem in terms of AI technologies.

There is a massive amount of data. For example, in a typical division there are almost 3000 units and subunits, each with separate identities, locations, equipment. At any one time, we may need to keep track of and reason about a large proportion of them.

There is a large degree of uncertainty inherent in any data received. This uncertainty may arise from enemy attempts at deception, ambiguity or a lack of detail in an observation, or information that is incorrect. The system will therefore require a capability for what we might call confirmatory reasoning, where no fact is accepted at face value, but with the reasoning system always looking for some confirming evidence.

The system will require fuzzy spatial reasoning, because Intelligence Officers use imprecise descriptions of position, such as "in the vicinity of", "on the left/right flank", "the boundaries of a Forming Up Place".

The system will require truth maintenance, that is, the ability to reuse parts of a chain of reasoning, when a fact or assumption is found to be incorrect(ref.7). It will also require what is called hypothetical reasoning, where it can reason about several competing explanations, or hypotheses, at the same time(ref.8).

Several clusters of knowledge, for matters such as doctrine, tactics, activities, terrain and equipment have been identified. It is anticipated that more clusters will be found as more domain experience is acquired. This partitioning suggests that a blackboard paradigm(ref.11) would be appropriate.

4. METHODOLOGY

The Intelligence Assessment problem, as has been described, is immense. However, a study of the domain has shown that the intelligence process reduces a large number of possible courses open to an opposing force to a manageable few likely ones(ref.6). This realisation has given us confidence that a computer-based decision aid is feasible, even if its implementation may be exceedingly challenging.

In this section, we describe the methodology we have adopted for this work. We discuss our work in knowledge acquisition and the results so far obtained. We describe our approach to system development and to how the developing system will be tested.

4.1 Knowledge acquisition

Our first move to acquire knowledge was to send two people to a military tactical intelligence course at the Australian Army School of Military Intelligence at Canungra. This course provided a basic understanding of the processes used in assessing information, and formed a starting point for the project.

As an aid to training, the Intelligence Corps has compiled a fictitious enemy whose characteristics are defined in a pamphlet titled "The Musorian Armed Forces". The characteristics of a particular enemy will be unique, but in the absence of an identifiable enemy, the MAF is the only one we have. The information in this book forms the basis of our knowledge of the way the enemy intends to organise itself and to fight.

We have an ex-Intelligence Officer working on this task. His primary task has been to extract, from the pamphlet, rules and facts which would be suitable for incorporating into an expert system of some kind. We have currently extracted somewhere over 1000 such pieces of doctrinal knowledge.

Our methodology has also relied on unstructured interviewing techniques and rapid prototyping. With the same ex-Intelligence Officer, we have worked through a standard exercise from the School of Military Intelligence, getting him to verbalise his knowledge. We have done some prototyping with this information, using the NEXPERT OBJECT expert system shell.

From this work, we have identified the potential knowledge sources cited in Section 3.2.5. We have enough detailed knowledge to have written a specification for an equipment knowledge source. Together with the Centre for Spatial Information Systems at CSIRO, we have also written a specification for a prototype inference engine (see Appendix I). Both these systems are scheduled for implementation between October 1990 and September 1991.

4.2 System development

Because of the difficulty inherent in the intelligence assessment problem, a key decision in terms of system development has been to "start small". We are taking a bottom-up approach. We are implementing those parts of a decision aid that we believe we understand, and hope, through prototyping and experimentation, to improve our understanding of the remainder of the system.

We have defined a base system (the CSIRO SPOT system), to be delivered in October 1990, which will be our phase one prototype. It should provide sufficient facilities to allow an Intelligence Officer to perform the same tasks that he does now. Our aim is to create a platform for the visualisation process and then add AI components.

We have partitioned the AI problem into knowledge sources. Using rapid prototyping techniques, we will prototype individual knowledge sources to enhance the functionality available in the phase one prototype. When sufficient knowledge sources have been prototyped, we shall be in a position to develop strategies to control the manner with which they interact.

After examining the commercially available knowledge representation toolkits and expert system shells against the criteria outlined in Section 3.2.5, the conclusion was that none of these AI tools was close to meeting our requirements. Had we found a suitable tool, there would still have been some effort required to build an interface to SPOT and its map management facilities. We therefore decided, together with CSIRO, to prototype a series of inference engines that will interface with SPOT. We plan to develop our inference engine prototype to match our understanding of the knowledge system as it is refined, with the prototyping of the AI components.

4.3 Testing

We intend to use an army training exercise, obtained from the School of Military Intelligence, as the basis of our prototyping. In Section 4.1, this exercise was described as the focus of a large part of our knowledge acquisition. We plan to use the exercise as a base for defining and developing the system, defining user's requirements and for testing the system as it is developed.

Our ex-Intelligence Officer is able to work through the exercise manually and obtain the expected results. The plan is to run progressively more of the same exercise on the computer, using the results of the manual exercise as a check.

5. DEVELOPMENT PLAN

5.1 Phase one

As has been stated, the SPOT system, developed by the CSIRO Centre for Spatial Information Systems, will be the the main output from phase one. The SPOT system will provide object data and manipulation interfaces to an ONTOS object database. A major component is an X-based Map Management Module. An Import Utility will convert digital terrain data in DLG3 format into topological constructs that can be stored in the object database. All these modules will be accessible through a Graphical User Interface. The SPOT system is shown in figure 2.

The Map Management Module allows a user to define terrain objects as objects of military significance, for example, observation posts or avenues of attack. A user can also create an ORBAT representing the opposing force. These objects, stored in the ONTOS database, can be displayed on map backgrounds to display a representation of the tactical situation. A user can also display any cultural, drainage or vegetation data, such as parts of a river system or small settlements. These facilities should provide a useful aid to an Intelligence Officer.

In addition, we have jointly developed with CSIRO a specification for an inference engine based on our understanding of the demands that the domain would make on an AI system. The specification appears as Annex A, and is discussed in the following section.

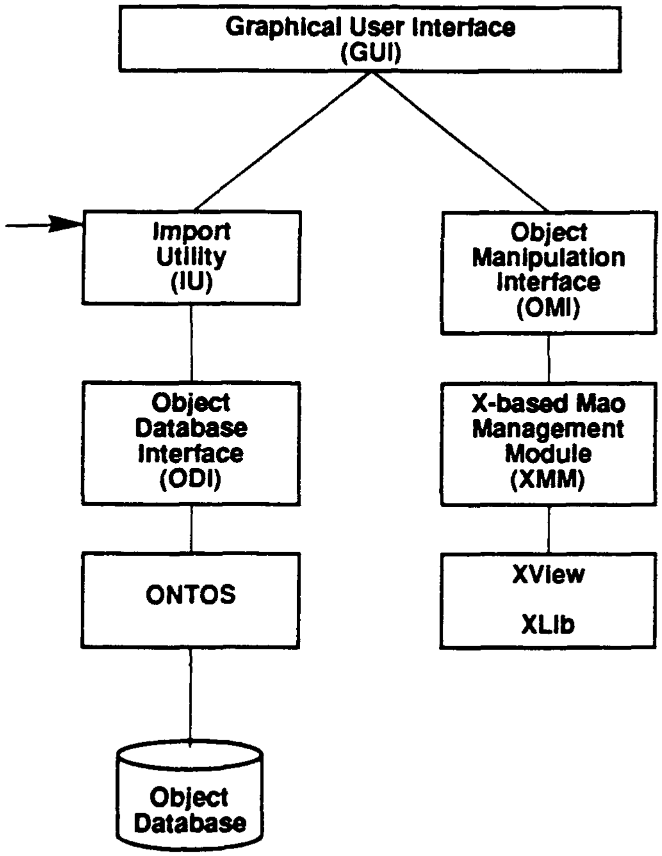

Figure 2. The SPatial Object Toolkit (SPOT)

5.2 Phase two requirements

In Phase 2, planned for between October 1990 and September 1991, work will continue to add and enhance the functionalities contained in the phase one prototype. We aim to implement a hypertext system. We will also implement a Message Input Module, which will record and parse new messages, and create any objects necessary for the inference engine.

We will develop our first prototype of our inference engine. We will prototype and develop functionality in terms of knowledge sources that will be specified. We have written a first specification for an Equipment Knowledge Source, and it is hoped a specification for a Doctrine Knowledge Source will follow shortly.

After several knowledge sources have been prototyped, it will be possible to investigate the interaction between knowledge sources and to develop models of cooperation.

With the delivery of our first prototype, we will be able to carry out a more formal exercise to elicit further user requirements.

On the basis of our experience with our prototype of the inference engine, and our improved understanding of the Intelligence Assessment problem, we will generate a specification for a further inference engine prototype. The first inference engine prototype will be rule-based.

Part of the work in phase two will be to investigate the applicability of expectation based reasoning to the problem. We hope to establish clear enough semantics for an expectation based reasoning system to write a specification for such a system during this phase.

5.3 Beyond Phase two

During phase three and beyond, we will continue to define user requirements, and develop the system accordingly. We will also refine the inferencing engine and prototype additional knowledge sources.

We recognise a need to apply techniques to different "sample" exercises. The exercises from the School of Military Intelligence, for reasons of instruction, are straightforward and lack ambiguity. To exercise our system's confirmatory reasoning capabilities, we would need more realistic exercises to test our systems.

The last Defence White Paper concluded that a low level intensity conflict is the most likely scenario Australia will have to face in the future. However, the scenario that our sample exercise assumes involves a conflict of medium intensity, because the doctrine for mid-level intensity conflicts is most well known and best documented. The system will be adapted as doctrine becomes available for low level intensity conflicts.

6. DESCRIPTION OF AI ASPECTS

6.1 Phase two work

It is planned, in phase two of this project, to incorporate several elements of artificial intelligence. We plan to implement a prototype inference engine and use it to begin to implement an Equipment Knowledge Source. On the basis of this work, we will specify a second phase of our inference engine. We will also investigate the semantics of expectation based reasoning system for the longer term.

6.1.1 *The equipment knowledge source*

We have chosen to first implement the Equipment knowledge source (EKS) because, as a result of our knowledge elicitation, it is the best understood of the knowledge sources. It also seems the easiest knowledge source to implement from an AI viewpoint.

The identification of equipment on a battlefield can indicate the composition of a force. The role of the EKS is to generate a list of unit-types, either individually or in groupings, that explain the equipment listed in a message. By unit-type, we mean the type and size of a unit, eg, a tank regiment or an artillery company. The EKS generally will not precisely identify units; the system will usually require more knowledge than is contained in the EKS for that task.

When implementing the EKS, the aim will not be to closely mimic the human reasoning process, but to achieve results comparable to that of an experienced Intelligence Officer.

The core knowledge of the EKS will be associations between units and the equipment they hold. It will also need knowledge regarding equipment doctrine. This is knowledge of which groupings of units are plausible, and how these groupings vary under changing circumstances, such as, for example, if the forces are advancing.

There is a class of equipment, called signature equipment, which uniquely identifies a unit type. For example, a T-62 tank would always indicate elements of a tank regiment, as they are the only units which hold this equipment. When a signature equipment is mentioned in an equipment list, the EKS will always focus its reasoning on that equipment first.

Two assumptions, used by Intelligence Officers, will underlie the working of the EKS. The first is that not all of the equipment in an area may be seen and reported in a message. The second is that the EKS will nominate the largest unit which might explain a sighting, in preference to a list of component units.

The Equipment knowledge source will propose hypotheses to explain equipment sightings, using cues such as equipment type and quantity, and units known to be in the area of the sighting. An hypothesis will suggest a unit-type or combination of unit-types to explain a sighting. In its first pass, it will generate all possible hypotheses up to some numerical limit.

When a sighting contains no signature equipment, rather than detailing all those units which might hold a particular item, the EKS will try to find a suitable level of abstraction to explain the sighting. For example, if a sighting included BTR-50 Armoured Personal Carrier, the hypotheses would be that these were elements of a Motor Rifle division, rather than generating many hypotheses which together detail all the units that make up a Motor Rifle division.

As each hypothesis is generated, the knowledge source will test whether all the equipment sighted has been accounted for. It will determine which equipment has not been accounted for. The most likely hypotheses will be those that best account for the sighted equipment. The knowledge source will try to merge hypotheses or add units to account for the unexplained equipment, using knowledge of plausible combinations of units and equipments.

If no satisfactory hypothesis can be generated, the EKS will postulate the presence of previously unsighted units.

Knowledge of doctrinal and location constraints can be used to prune the search. For example, if the reported unit is stationary, it is worth checking whether any units are known to be in that location and if they could explain the equipment seen. As another example, if the location of a sighting falls within where the second echelon of a force is assumed to be, doctrine will focus the EKS on second echelon units as candidate solutions.

Unless a user has nominated a hypothesis as being correct, the output will be all those hypotheses that explain most or all of the equipment sightings and satisfy some sort of plausibility criterion.

6.1.2 *Visualisation, explanation and user interaction*

It is not intended that the system will present one hypothesis as being correct, and discard the remainder. It will present all hypotheses to the user for consideration. He can examine the tactical implication of each hypothesis by selecting it for display by the Map Management Module. It will superimpose the iconic representations of the hypothesised units on the map background.

At any point in the reasoning process, a user will be able to nominate a hypothesis as being the best to proceed with or to be discarded. Alternatively a user may wish to suggest a different hypothesis.

We have a strong preference for producing multiple competing hypotheses rather than using numeric approaches to uncertain reasoning. From our reading of the literature, numeric approaches do not seem to provide adequate explanations as to the causes of uncertainty in a conclusion(ref.10). It is also not clear how to mix the semantics of probabilistic reasoning and truth maintenance.

We are attempting to assist a user in making a better decision. We feel that, given the difficulties in characterising the domain and deriving a priori probabilities, confidence factors would provide a false degree of precision within conclusions. Our aim is to make all the supports for a conclusion explicit, and to allow the user to judge their validity.

In this first prototype of an inference engine, explanation will involve the standard method of supplying a trace of rules fired. If the hypertext system can be implemented during this phase of the work, it is proposed that there will be a link between all rules and the text in the Musorian pamphlet from which the rule was extracted.

6.1.3 *First inference engine specification*

The first inference engine specification is based on the characteristics of the intelligence assessment problem as indicated by the knowledge acquisition phase of the project. In particular, it has been shaped by the requirements imposed by the specification of the Equipment knowledge source.

It has already been stated that the first version of the inference engine is a prototype. As the inference engine and the Equipment knowledge source are developed, and as further knowledge is acquired with respect to other knowledge sources, we will refine our specifications.

The document included in Annex A is the current version of a working paper being written jointly by DSTO and CSIRO. Given resource constraints, it will be possible to implement only part of the specification. Those aspects, if any, which will be omitted are currently under discussion.

6.2 Expectation based reasoning

A significant portion of the behaviour of an opposing force will often correspond to some predefined behaviour. For example, before a deliberate attack, the units involved will move into a Forming Up Place. There may be radio silence in the hours preceding an attack.

Analysis already performed for the Australian Army Tactical Command and Control Testbed project has identified event and doctrinal templates(ref.4). We plan to identify temporal, spatial and doctrinal templates, and use them to confirm intelligence reports and to anticipate the actions of the opposing force.

Reasoning would become a matter of seeking cues indicating that a particular template may or may not be active. We would look for reports of events that might confirm or refute the expectations that a template would create. The system would need sufficient flexibility to allow for an amount of variation in the timing or location of events for them to fit within the expected patterns.

We anticipate that an expectation based reasoning system would be built around a frame-based reasoning system. There would need to be mechanisms to express expectations and to associate them with their parent templates. The system would need mechanisms to monitor for the presence or absence of these expectations.

The system would need to associate strengths between events and expectations. For example, a particular event may be thought of as vital to confirm a particular template. Its absence within a specified period may be taken as refutory evidence. Some other event may only weakly confirm a particular template.

We plan to investigate the semantics of such a system during phase two. If a system can be specified, we hope to implement it in subsequent phases of the work.

7. SUMMARY

This document has described the work on a project to develop a decision aid to assist an Intelligence Officer in a tactical HQ with the interpretation of intelligence concerning opposing force activities. The paper has described the framework of the project, in particular, the need to merge several diverse technologies, to add artificial intelligence features to augment existing systems and to adopt a human-centred approach to the interaction between user and machine.

The work in eliciting knowledge from an Intelligence expert, and a collaboration with a division of CSIRO to develop a map-based spatial reasoning system have been described. We have also described our plans to develop an inference engine with support for hypothetical reasoning and truth maintenance, and to prototype one or more knowledge sources in the next twelve months.

REFERENCES

No.	Author	Title
1	Andrews, T., Harris, C. and Kiril Sinkel	"The ONTOS Object Database". Ontologic Inc., 1989
2	Andriole, S.J. and Hopple, G.W.	"Defense Applications of Artificial Intelligence". Lexington Books, 1988
3	Atkinson, M. et al	"The Object-Oriented Database System Manifesto". Rapport Technique Altair 30-89, 21 August 1989
4	AUSTACCS Project	"AUSTACCS Tactical Command Support System Specification". AUSTACCS Project Office, Army Office, Canberra, 1990
5	Ben-Bassat, M.	"Knowledge Requirements and Management in Expert Decision Support Systems for (Military) Situation Assessment". IEEE Trans. Systems, Man and Cybernetics, Vol SMC-12, No. 4, 1982
6	Calder, P.F.	"The Importance of Terrain in Military Intelligence". Presented to Seminar "Geographic Information in the Defence of Australia" June 1988
7	de Kleer, J.	"An Assumption-based TMS". Artificial Intelligence, Volume 28 (1), 1986
8	Filman, R.E.	"Reasoning with Worlds and Truth Maintenance in a Knowledge-Based Programming Environment". Communications of the ACM, Volume 31 Number 4, April 1988
9	Hollnagel, E., Mancini, G. and Woods, D.D. (eds)	"Cognitive Engineering in Complex Dynamic Worlds". Academic Press, 1988
10	Luger, G.F. and Stubblefield, W.A.	"Artificial Intelligence and the Design of Expert Systems". Benjamin / Cummings, 1989

11	Nii, H. Penny	"Blackboard Systems". Report No. STAN-CS-86-1123, Dept. of Computer Science, Stanford University, June 1986
12	Roth, E.M., Bennett, K.B. and Woods, D.D.	"Human Interaction with an "intelligent" machine". (Hollnagel 1988)
13	Spain, D.S.	"Application of Artificial Intelligence to Tactical Situation Assessment". IEEE EASCOM, 1983
14	Smith, J.	CSIRO Centre for Spatial Information Systems, Canberra. Personal communication
15	Taylor, E.C.	"Artificial Intelligence in Command and Control - What and When?". (Andriole 1988)
16	Woods, D.D.	"Commentary: Cognitive Engineering in Complex Dynamic Worlds". (Hollnagel 1988)

ANNEX A

SPECIFICATION FOR FIRST PROTOTYPE OF INFERENCE ENGINE

The following functionality is required in the initial experimental platform:

1. Rule Inferencing and Rule Classes

This type of kernel is specified at this stage because of the maturity of this technology. In addition the implementation should be open ended so that additional functionality (ultimately, schema based reasoning) can be provided. For example, functions implemented at the lower level can manipulate the knowledge base state space and be used as elements in the rule language. This mechanism should be recursive. In other words, well formed constructs in the underlying language can be evaluated as conditions or actions within a rule.

Rule classes will allow the partitioning of rules in a knowledge base. All rules must be members of a rule class, and the rule superclass. It should be possible to form hierarchies of rule classes.

Each rule will inherit from the rule superclass a graphics/iconic object, which can be used for debugging.

Rules can be one of two types: same world or new world.

Rules can be either forward or backward chaining. The syntax of a rule will be the same for both types. The direction in which a rule (more precisely, the rules in a rule class) will run will be determined at run time.

There should be debugging facilities built into the rule system. It should at least be possible to set breakpoints (in antecedents and consequents), and to single step through the rules.

One must be able to choose between conflict resolution strategies.

It should be possible to at least view, and preferably have, direct access to and manipulation of the system agenda.

2. Worlds Mechanism

This is essential in all segments of the message inferencing problem.

3. Tagging

An essential capability of the inferencing mechanism will be to handle the non-monotonic nature of the reasoning process. Antecedents for any state will have to be accessible to a truth maintenance process, to handle situations in which a message (or its associated inferences) contradicts the inferences drawn from previous messages.

An implicit requirement of the decision support model is explanation. This also requires a tagging mechanism, so that conclusions can be related to the information which triggered them.

4. Visualisation

A requirement is for visualisation in the two dimensional mapping context. Message processing will generate dynamic objects, some with non-monotonic characteristics, for each world.

Another requirement is for a world window providing access to the knowledge base and state space of each world, and differential comparisons between worlds.

Mouse, menu and language interfaces to knowledge bases are required.

The implementation of an ad-hoc query facility with some degree of pattern matching (aka KEE's TELL AND ASK system) is desirable.

5. User Interaction

The visualisation channels should be input/output. The user will provide the high level control of the inferencing process and in some cases will force pruning of the multiple worlds. Often this will be done through assertions in a process akin to knowledge acquisition. These assertions should be recorded by the tagging function.

6. Object Database Interface

A global spatial object database will serve many modules in a complex decision support system. Each inferencing mechanism will require a database interface which supports predicate evaluation, iterative instantiation of variables, property and relationship evaluation.

7. Object System

The object system should be used to integrate the various parts of the inferencing system. Therefore, each part of the inferencing system should ideally be modular, to allow any part that is extraneous for a particular application to be stripped out for efficiency. To facilitate the integration process, rules and rule classes should be objects.

The object system should support value class checks, eg, should be able to specify that a slot value is a member of a designated class.

The system should support access oriented programming, allowing methods to be activated automatically when a specified slot is accessed (aka KEE's Active Values).

8. Temporal Reasoning

There should be support for some simple temporal reasoning mechanism. As a minimum, all events should be time-stamped, and it should be possible to perform standard arithmetic operations on the time values.

In later versions, temporal reasoning can be either interval or point based. The system should handle fuzzy temporal concepts such as past, present and future.

9. Language Interfaces

There should be interfaces to other languages, eg C++, LISP, PROLOG.

Modules should be able to run external programs and use results, and be invoked by external programs and return results.

10. Graphical Interface

A system to allow graphical display of selected slots in selected objects (aka KEE's Active Images) is required.

11. Debugging and Tracing

The system should supply a range of debugging aids, allowing a developer to trace, monitor, interrupt, interrogate and modify the running of the system.

For the rule system, there should be:

- a stepper option to allow rules to be fired one at a time,
- a break option on each rule (in antecedents and consequents),
- ways of structuring rule traces where large number of rules have been fired,
- a debugger at a higher level than the underlying language debugger.

12. Relationship Definition System

There will be a kernel of objects defining classes of relationships and their properties. It will allow the definition of relationships and their inverses, ideally graphically, and with the system performing the housekeeping. It will allow the specification of classes of relationship, eg, whether a relationship is transitive. Ideally, normal inheritance relationships, such as is-a and instance-of, should be implemented as part of this system.

13. Architecture

There should be a common data area to allow knowledge sources and applications to communicate.

DISTRIBUTION

	No. of copies
Defence Science and Technology Organisation	
Chief Defence Scientist First Assistant Secretary Science Policy First Assistant Secretary Science Resources Planning and Commercialisation Director General Science and Technology Programs Director General Space and Scientific Assessments Assistant Secretary Science Corporate Management Assistant Secretary Development Projects	1
Scientific Adviser to Defence Central	1
Counsellor, Defence Science, London	Cnt Sht Only
Counsellor, Defence Science, Washington	Cnt Sht Only
Defence Science Representative, Bangkok	Cnt Sht Only
Scientific Adviser, Defence Research Centre, Kuala Lumpur	Cnt Sht Only
Weapons Systems Research Laboratory	
Director, Weapons Systems Research Laboratory	1
Chief, Combat Systems Division	1
Research Leader, Combat Systems	1
Head, Combat Systems Technology	1
Head, Tactical Command Information Systems	1
Authors	2
Electronics Research Laboratory	
Director, Electronics Research Laboratory	1
Chief, Information Technology Division	1

Head, Information Systems Research Group	1

Libraries and Information Services

Librarian, Technical Reports Centre, Defence Central Library, Campbell Park	1
Manager, Document Exchange Centre, Defence Information Services	1
United States, Defense Technical Information Center	2
United Kingdom, Defence Research Information Centre	2
Canada, Director Scientific Information Services	1
New Zealand, Ministry of Defence	1
National Library of Australia	1
Main Library, Defence Science and Technology Organisation Salisbury	2
Library, DSD, Melbourne	1
Library DSTO Sydney	1
Library MRL	1
Library EDE	1
Australian Defence Force Academy Library	1
British Library, Document Supply Centre	1

The Technical Cooperation Program (TTCP)

Australian National Leader, Subgroup S	1
Australian National Leader, Subgroup X	1

Department of Defence

AGPS	1
Scientific Adviser to Defence Intelligence Organisation	1

Army Office	
Scientific Adviser - Army	1
Director, Military Intelligence	1
School of Military Intelligence, Canungra	1
CSIRO	
Division of Information Technology, Canberra (Attention: Dr J. Smith)	1
Spares	4
Total number of copies	41

DOCUMENT CONTROL DATA SHEET

Security classification of this page : UNCLASSIFIED

1 DOCUMENT NUMBERS	**2 SECURITY CLASSIFICATION**
AR Number : AR-006-707	a. Complete Document : Unclassified
Series Number : WSRL-TM-18/91	b. Title in Isolation : Unclassified
Other Numbers :	c. Summary in Isolation : Unclassified
	3 DOWNGRADING / DELIMITING INSTRUCTIONS

4 TITLE

A DECISION SUPPORT SYSTEM FOR MILITARY INTELLIGENCE ASSESSMENT

5 PERSONAL AUTHOR (S)	**6 DOCUMENT DATE**
N.G. Gori and P.F. Calder	March 1991
	7 7. 1 TOTAL NUMBER OF PAGES 20
	7. 2 NUMBER OF REFERENCES 16
8 8. 1 CORPORATE AUTHOR (S)	**9 REFERENCE NUMBERS**
Weapons Systems Research Laboratory	a. Task : 779153
	b. Sponsoring Agency :
8. 2 DOCUMENT SERIES and NUMBER Technical Memorandum 18/91	**10 COST CODE**
11 IMPRINT (Publishing organisation)	**12 COMPUTER PROGRAM (S)** (Title (s) and language (s))
Defence Science and Technology Organisation	

13 RELEASE LIMITATIONS (of the document)

Approved for Public Release.

Security classification of this page : UNCLASSIFIED

14 ANNOUNCEMENT LIMITATIONS (of the information on these pages)

No limitations

15 DESCRIPTORS		16 COSATI CODES
a. EJC Thesaurus Terms	Artificial intelligence Military intelligence Information systems Command and control	120901, 1504
b. Non - Thesaurus Terms		

17 SUMMARY OR ABSTRACT
(if this is security classified, the announcement of this report will be similarly classified)

(U) This paper details the early stages of research towards designing a decision aid to assist an Intelligence Officer in a tactical HQ. The particular problem being tackled, the interpretation of intelligence messages concerning opposing force activities, is described, as are those factors contributing to the task's difficulty.